REDUCING THE VARIANCE OF CYCLE TIMES IN SEMICONDUCTOR MANUFACTURING SYSTEMS

M. Mittler, A.K. Schömig, and N. Gerlich
Bayerische Julius-Maximilians-Universität Würzburg,
Lehrstuhl für Informatik III
Am Hubland, 97074 Würzburg, Fed. Rep. of Germany
Tel.: +49-931-8885518, Fax: +49-931-8884601
e-mail: [mittler|schoemig|gerlich]@informatik.uni-wuerzburg.de

Abstract: *In semiconductor manufacturing, due to rework and re-entrant flow, overtaking of wafers can occur. The effect of overtaking is that cycle times at successive service centers are not mutually independent. As far as the distribution of cycle times is concerned, only higher moments are affected, the mean cycle time remaining unchanged by the influence of overtaking. Further, in the literature, it is conjectured that variance of cycle times increases when overtaking increases.*

Taking into account this conjecture, we attempt at reducing the variability of cycle times by diminishing the magnitude of overtaking. This can be done by reversing the overtaking through appropriate sequencing rules. In order to achieve this goal, we examine several sequencing rules by means of simulation studies based on real data sampled at four different semiconductor manufacturing facilities. Our results elucidate that there is no general correlation between the magnitude of overtaking and the variance of cycle times. We show that the performance measures under consideration depend very much on the type of the facility and the product mix.
Keywords: *Semiconductor Manufacturing, Scheduling, Cycle Times, Simulation*

1 Introduction

Semiconductor manufacturing is among the most complex manufacturing processes. A semiconductor chip is a highly miniaturized, integrated electronic circuit consisting of thousands of components. Depending on the scale of integration, the type of chip, and customer specs, the whole manufacturing process may require up to 500 single steps. Most of these operations involve cleaning, deposition, lithography, etching, ion implantation and testing. Note, there are no assembly operations before reaching the backend stage of the production process. Several performance measures are used to assess a semiconductor manufacturing facility. To highlight the most important, we mention machine utilization, machine downtimes, production yield, throughput, and cycle time. We define cycle time as the time a lot of wafers needs to travel from start to finish the core manufacturing process. The capability of meeting due dates has become a crucial factor in global manufacturing competitiveness. Thus, a modern manufacturing organization has efficiently to deal with this particular issue in order to be able to compete on the global marketplace. Consequently, not only the mean cycle times but also the variance of cycle times have to be minimized.

Standard literature on scheduling (e.g. Baker [2], Pinedo [10]) focuses on average cycle times and system utilization only. Many studies deal with the reduction of mean cycle times by setup avoidance strategies, sophisticated control policies to form batches (cf.Gold and Frötschl [4], and Tran-Gia et al. [12]) or the provision of redundant tools (cf. Leachman et al. [7] and Xu et al. [15]). Recent studies devote attention also to the reduction of the variability of cycle times. For example, Kumar [5] investigates how to achieve this goal by means of appropriate sequencing rules.

In this work we focus on the problem of rework and reentrant the flow of wafers. Due to these factors, overtaking of wafers can occur. There are several categories of overtaking. In the case of physical overtaking, a wafer that entered a particular processing step or the entire manufacturing system ahead of another wafer is finished after the second wafer. Overtaking does not necessarily have to be physical; it is only required that the probabilistic effects caused by a lot during its visit at a particular tool propagate through the manufacturing facility, so as to affect its cycle time at subsequent tools. The effect of overtaking is that cycle times at successive service centers are not mutually independent if the process flow is not overtake-free. For detailed discussions of overtaking we refer to Melamed [8], Mittler and Gerlich [9], and Walrand and Varaiya [13]. As far as the distribution of cycle times is concerned, only higher moments are affected, the mean cycle time remains unchanged by overtaking. Further, Whitt [14] states that variance of cycle times increases when overtaking increases. To treat the phenomenon of overtaking quantitatively Whitt defined a measure for the magnitude of overtaking. This measure is referred to as the *amount of overtaking* and is defined as the mean number of customers $\mathrm{E}[N]$ overtaken by an arbitrary customer (*active view*), or the mean number of customers $\mathrm{E}[N^\star]$ an arbitrary customer is overtaken by (*passive view*).

Taking into account the aforementioned conjecture, we attempt at reducing the variability of cycle times by diminishing the magnitude of overtaking. This can be done by reversing the overtaking through appropriate sequencing rules. We investigate several sequencing rules in order to achieve the goal of reduced variance of cycle times. We evaluate the performance of several rules by means of simulation based on specifications of real semiconductor manufacturing systems.

2 Research Methodology

2.1 Experimental Design

According to our research objectives, we selected two experimental factors: (1) sequencing rule, and (2) capacity utilization. Implicitly we have to deal with a third factor: type of product.

Sequencing rule: For we attempt at reducing the variance of cycle times by diminishing the overtaking we examine the data set shown in Tab. 1 for the following sequencing rules: *First-In First-Out* (FIFO), *Earliest Due Date* (EDD), *Critical Ratio* (CR), and *Closest to Completion by Step* (CC-Step).

Since FIFO is widely used in manufacturing we compare the results of all sequencing rules to FIFO. Due to the mechanics of batch formation of jobs is is possible that a late–arriving job can bypass earlier jobs of another job type to become part of a batch. There are also cases when operators are needed for service at a tool, where a job cannot be processed since there is a special operator for this particular job type missing. This job might be overtaken by a job of a different type that can be serviced immediately. Note, that these scenarios can happen without specifying priorities. Other sequencing rules are overruled by batch forming laws as well.

EDD is taken into account because the maximum gain that can be achieved by reversing the overtaking is obtained by sequencing the customers in the order of their first arrival to the semiconductor manufacturing facility. EDD is applied with no due dates set explicitly such that a lot receives service prior to another if it has a longer system lifetime. By this policy, the order of the customer arrival is preserved as much as possible.

Additionally, we apply CR and CC-Step. CC-Step selects the job with highest ratio of the current step number and the total number of processing steps in the job's process flow for service to be given next whereas CR, in case of equal priorities, favors the job with the lowest critical ratio c which is based on total remaining processing time ($TRPT$), the due date (DD), and the current time (NOW):

$$c = \begin{cases} \dfrac{1 + DD - NOW}{TRPT}, & DD > NOW, \\ \dfrac{1}{(1 + NOW - DD) \cdot TRPT}, & DD \leq NOW. \end{cases}$$

For further details, we refer to the original documentation [3] of the Delphi simulation tool. The reason for the application of CR and CC-Step is that these sequencing rules also provide a means to distinguish at a particular tool among the lots of the same type waiting for different processing steps to be carried out. These rules favor either those lots that wait for a processing step located back in the job's process flow or those lots that were already reworked.

Capacity utilization: We introduced this factor to determine whether the results depend on the load at which the facility is working. The notion of *capacity* is defined by the Delphi tool as the maximum production rate of the *expected bottleneck tool*. Delphi automatically tries to find this bottleneck tool in a preprocessing run using queueing analysis. Since the utilization of any tool depends on factors such as rework, scrap, operator unavailability, and product mix that can only be approximated by a queueing model it is possible that the *actual* bottleneck is different from the expected bottleneck tool. We examined the capacity utilization at two levels: a low level of 30 percent and a high level of 90 percent. This choice of 90 percent is justified since we expect that most semiconductor manufacturing facilities are more or less heavy loaded.

Product type: We simulate semiconductor fabrication facilities using factory–level *data sets* of four different semiconductor fabs. These data sets were collected during the MIMAC project. These data sets include ASICs, microprocessors, and

non-volatile memory, and can be obtained via anonymous FTP from SEMATECH. An overview of these data sets is given in Tab. 1. We keep the original denotation of the data sets to avoid any misunderstandings.

Each wafer type requires a certain sequence of processing steps, maybe receiving service from diverse tools. Note, that the product mix is given by the portion of one wafer type of the overall load. As a consequence, the dispatching rate of lots of wafers may differ significantly. The process flow specifications for each jobtype are given separately. Hence, we expect different results for every wafer type.

If not mentioned explicitly, we set all other parameters, like rework, machine breakdowns, setup etc., according to the original data sets.

2.2 Simulation Parameters

This study was conducted using the Delphi simulation tool, a package for simulating large queueing networks, with an emphasis on providing the building blocks for semiconductor manufacturing simulation. Delphi was developed by Frank Chance and Lee Schruben at Cornell University, Ithaca, U.S.A., and then tested and improved by the Performance Modelling Group of Universität Würzburg, Germany. There is no charge for Delphi if it is used for research at a degree granting institution. Moreover, we extended Delphi for the purpose of simulating the amount of overtaking.

We use a full factorial design of the first two factors. The third factor is implicitly included, since each simulation run incorporates all wafer types. The entire design consists of 8 design points. For each design point we run a simulation of 50,000 hours simulated time on a SUN SPARC workstation. We discarded statistical data from the initial 10,000 hours to avoid starting condition effects. Delphi automatically conducts a statistical test procedure (cf. Schruben [11]) to check for transient effects in the simulation output. This is done by calculating a test statistic that is sensitive to changes in the batch means. The p–value of this statistic is the lowest significance level at which one would reject the hypothesis, that there is no transient effect in the simulation output. According to Chance's recommendation [3], p.44, we choose the simulation time such that a reasonable amount of CPU time had to be used and the Schruben test did not indicate nonstationary simulation output.

Table 1: *SEMATECH Testbed Data Sets*

Data Set	1 : Logic_1	3 : Logic_2	5a : ASIC_1	6 : ASIC_2
Type of Factory	Commodity	Commodity	ASIC	ASIC
Product Mix	2	11	21	9
approx. WSPM	16000	21400	10000	5500
Avg. #Mask Layers	15	35	30	30
#Processing Steps	210 / 245	298 – 533	117 – 259	234 – 356
#Tools	83	73	85	104

2.3 Analysis of Results

Since scrapping is possible, one cannot use Little's law to estimate the cycle time for all jobs, because scrapped jobs are included despite of the usual notion of cycle time. Hence, Delphi uses transaction observations of those jobs that exit normally to estimate cycle times. By default Delphi splits the output data into 20 batches. Batch means are calculated and under the assumption of independence and normal distribution of the batch means an estimate for the mean and the variance is derived. These batch means are also used to calculate the test statistic mentioned above. Confidence intervals are derived in straightforward manner.

To calculate the amount of overtaking we added the following procedure: Every time a job leaves the fab the number of jobs that arrived before this job and still are in the fab is calculated. Each of this jobs has a counter for the number of jobs it is overtaken by that is incremented by 1. Such samples for the active and passive amount of overtaking are provided. For each product type the active and the passive amount of overtaking is calculated from samples. The overall amount of overtaking is obtained by a weighted summation of the individual product amounts:
$\mathrm{E}[N] = \sum_i \frac{\lambda_i}{\lambda}\mathrm{E}[N_i]$, $\mathrm{E}[N^\star] = \sum_i \frac{\lambda_i}{\lambda}\mathrm{E}[N_i^\star]$, where λ_i is the arrival rate for product i and $\lambda = \sum_i \lambda_i$.

We analyzed the simulation data in two steps. First, we compared the cycle times of each sequencing rule with FIFO using the *paired–t test* to check whether the sequencing rule affects the mean of the cycle times. Then, we looked at the amount of overtaking and conducted a *Snedecor–F test* to detect a significant difference in the variance of the cycle times. This was done for each product and capacity utilization.

The Paired–t Test

A paired–t test is carried out to check whether two system alternatives differ significantly in terms of a performance measure (cf. Law and Kelton [6]). The test is constructed such that the difference in the two performance measures will lead to the rejection of the hypothesis "the two system alternatives are the same".

Given are samples $X_{i1}, X_{i2}, \ldots, X_{in_i}$ of n_i independent identical distributed observations taken from system i, $i = 1, 2$. System 1 is commonly called *the standard*. In our case the X_{ij} are batch means and the number of batches is given as $n = n_1 = n_2 = 20$. Let $Z_j = X_{1j} - X_{2j}$ be the paired difference of observations (batch means) for $j = 1, 2, \ldots, n$. Then the Z_j's are independent identical distributed random variables with expected value $\mathrm{E}[Z_j]$.

We now can construct a confidence interval for $\mathrm{E}[Z_j]$ using $\overline{Z}$, the sample mean, and $\mathrm{Var}[\overline{Z}]$, the sample variance. Hence, if the Z_j's are normally distributed, the $100 \cdot (1-\alpha)$ percent confidence interval is given by $\overline{Z} \pm t_{n-1,1-\alpha/2}\sqrt{\mathrm{Var}[\overline{Z}]}$.
This interval is called the *paired–t confidence interval*. Note, that we do not have to assume that X_{1j} and X_{2j} are independent, nor that $\mathrm{Var}[X_{1j}] = \mathrm{Var}[X_{2j}]$.

The paired–t confidence interval overlaps with a probability of $100 \cdot (1-\alpha)$ percent the true value of the difference in the performance measure of the two system

alternatives. Therefore the hypothesis that the two systems under investigation are different has to be rejected, if the paired–t confidence interval does not include zero.

The Snedecor–F Test
We compare the variability of two samples drawn from different populations. Our hypothesis is that the two populations have the same variance, i.e. $\sigma_1^2 = \sigma_2^2$. We use an algorithm presented by Allen [1] to check for inequalities of the variances of cycle times.

Again let $X_{i1}, X_{i2}, \ldots, X_{in}$ be samples drawn from two populations connected with system i, $i = 1, 2$. Let s_i^2 be the sample variance. The ratio of the sample variance $f = \frac{s_1^2}{s_2^2}$ is Snedecor–F distributed with $\nu_1 = n_1 - 1$ and $\nu_2 = n_2 - 1$ degrees of freedom. The hypothesis $H_0 : \sigma_1^2 = \sigma_2^2$ has to be rejected if
$f > F_{\alpha/2,n_1-1,n_2-1}, \quad (s_1^2 > s_2^2) \quad \text{or} \quad f < F_{1-\alpha/2,n_1-1,n_2-1}, \quad (s_1^2 < s_2^2)$.
According to Allen [1] we exchange s_1^2 and s_2^2 such that $s_1^2 \geq s_2^2$ and use the alternative hypothesis $H_1(b) : s_1^2 > s_2^2$. In case of this alternative hypothesis the critical region for hypothesis H_0 is the set of all $f > F_{\alpha,n_1-1,n_2-1}$. For our parameters the corresponding value of the F distribution can be found as: $F_{0.1,19,19} = 1.82$.

Using this test we can determine, at the 10 percent level of significance, whether it is reasonable to suppose the variance in the cycle time for two systems is the same. However, we have to underline the fact that this test like all test on variances and unlike procedures for tests on the mean are not robust to the normality assumption.

3 Results

We now compare the results of all sequencing rules to FIFO. The tables show both, the paired-t confidence intervals for mean cycle times and the results of the two-sample test of variance of cycle times. To clarify the presentation we also depict the confidence interval of cycle times in the first line of the table showing the paired-t confidence intervals. The confidence level was set to 90 % in each case.

Data Set 1 The results obtained for data set 1 exhibit that by means of sophisticated queueing disciplines like EDD and CR the amount of overtaking can be reduced (see Table 2). Note, CR achieves less overtaking compared to EDD. The amount of overtaking from both active and passive point of view are almost the same. Whitt [14] showed that this relation generally holds for time-reversible queueing systems. However, the inversion of this argument is not permitted. Besides these observations the most important fact is that there is no significant difference among the sequencing rules in terms of the mean cycle time E[T] with exception of EDD for product 1 (see Table 3). We do not have a general explanation of this effect, but we believe that due to batch accumulation times longer waiting times might occur with employing EDD.

Remembering the conjecture of Whitt, i.e. the variability of cycle times increases with increasing overtaking, the results of our attempt at reducing the variability of cycle times by reducing the amount of overtaking are disappointing. The variance of

Table 2: *Active and passive amount of overtaking for data set 1*

Seq. Rule	E[N]	E[$N^\star$]
FIFO	31.03	31.36
EDD	14.86	15.02
CC-Step	31.67	32.03
CR	3.05	3.07

Table 3: *Paired–t confidence interval, data set 1*

Seq. Rule	lower	E[T]	upper	lower	E[T]	upper
FIFO	745.1	761.4	777.8	1010.1	1036.1	1062.1
		E[Z]			E[Z]	
EDD	-134.0	-115.1	-96.2	-3.2	18.7	40.5
CC-Step	-21.6	-3.0	15.6	-38.9	-6.8	25.3
CR	-282.9	-237.2	-191.5	-37.9	14.6	67.1
	product 1			product 2		

cycle times Var[T] does not vary significantly among the sequencing rules employed. For a significance level of $\alpha = 0.1$ the hypothesis that there is no difference between the variance obtained by FIFO and EDD, CC-Step, or CR, respectively has always been accepted with exception for sequencing rule CR and product 1, as shown in Table 4. Although the amount of overtaking is smaller for CR the corresponding variance of cycle time is much higher compared to FIFO. Thus, no improvement can be achieved by invoking sophisticated sequencing rules.

Data Set 3 We slightly modified the data set since the original one caused serious problems when running EDD at a load of 90 % of the bottleneck tool. When EDD was invoked the inventory of the fab turns to infinite with progressive simulation time. The reason for this is that at station 'FURN3' the inventory piled up at this tool. We overcame this problem by replacing sequencing rule EDD by FIFO at this station. For the purpose of statistical evaluation, we consider only products 1 to 6 for which a considerable number of lots have left the fab.

The amount of overtaking obtained for data set 3 is shown in Table 5. Once again, as for data set 1, employing a sequencing rule that makes use of due dates results in a smaller amount of overtaking. In this case, EDD outperforms CR in

Table 4: *Snedecor–F test of variance, data set 1*

Seq. Rule	Var[T]	f	H_0	Var[T]	f	H_0
FIFO	2900.57			6160.34		
EDD	3178.83	1.10	acc.	4403.21	1.40	acc.
CC-Step	2561.19	1.13	acc.	7556.01	1.23	acc.
CR	7686.53	2.65	rej.	6453.90	1.05	acc.
	product 1			product 2		

Table 5: *Active and passive amount of overtaking for data set 3*

Seq. Rule	E[N]	E[$N^\star$]
FIFO	41.72	41.73
EDD	27.02	27.02
CC-Step	42.01	42.01
CR	40.91	40.91

Table 6: *Paired–t confidence interval, data set 3*

Seq. Rule	lower	E[T]	upper	lower	E[T]	upper	lower	E[T]	upper
FIFO	533.7	538.5	543.4	354.1	356.6	359.0	621.9	627.7	633.6
		E[Z]			E[Z]			E[Z]	
EDD	-0.6	4.8	10.1	-73.6	-68.1	-62.5	18.5	24.6	30.7
CC-Step	-9.8	-4.5	0.9	-3.2	-1.0	1.2	-9.2	-2.5	4.2
CR	-321.6	-317.1	-312.6	-656.5	-652.7	-648.9	-138.4	-132.1	-125.9
	product 1			product 2			product 6		

terms of overtaking. Let us turn now to cycle times. Generally, sequencing rule CR achieves the worst mean cycle times. For products 1 to 5, among the remaining sequencing rules, there is either no significant difference or the difference is rather small (cf. Table 6 where only representative results are shown). CC-Step performs as well as FIFO for all products under consideration since the paired–t confidence interval contains zero. For product 6 only, EDD outperforms FIFO in terms of mean cycle times. However, this relation turns upside down in case of product 2.

Generally, the two-sample test of variance for a confidence level of 90 % exhibits that there is no significant difference among the variances of cycle times achieved by using sequencings rule EDD, CC-Step, and CR compared to FIFO. We noticed only few exceptions which are shown in Table 7. The results for EDD are never better than those obtained employing FIFO, even for products 2 and 3, EDD performs significantly worse. CC-Step performs as well as FIFO. Furthermore, CR significantly achieves a smaller variance for product 4 compared to FIFO but a larger variance for product 3.

We conclude the presentation of the results for data set 3 by discussing the modification of the original data at station 'FURN3'. Remember, that since the fab does not reach steady state at a load of 90 % of the expected bottleneck when

Table 7: *Snedecor–F test of variance, data set 3*

Seq. Rule	Var[T]	f	H_0	Var[T]	f	H_0	Var[T]	f	H_0
FIFO	254.97			339.62			862.11		
EDD	857.93	3.36	rej.	938.74	2.76	rej.	747.75	1.15	acc.
CC-Step	302.58	1.19	acc.	400.07	1.18	acc.	942.58	1.09	acc.
CR	170.94	1.49	acc.	1277.59	3.76	rej.	455.41	1.89	rej.
	product 2			product 3			product 4		

employing sequencing rule EDD we replaced EDD by FIFO. One might conjecture that this approach is misleading. However, we examined the *original* data of fab 3 for input percentages varying from 70 to 80. We also ran fab 3 with the modified data set. There is no general relation about whether invoking FIFO at station 'FURN3' leads to better cycle times. It decreases mean cycle times for some products as well as in some cases higher mean cycle times are achieved.

The results obtained for the modified data set are the same as for the original data set such that in terms of E[T], FIFO achieves almost the same results as CC-Step, FIFO performs better than CR, and the competition among FIFO and EDD is tied. Furthermore, the variance of cycle times also depends on the considered product; Table 7 gives the representative relations among the the four sequencing rules for the original data set. The variance achieved by CC-Step is likely to the one of FIFO which mostly performs better than EDD. Finally, only for a single product, invoking CR leads to a smaller variance compared to FIFO, otherwise the performance of FIFO is either better or comparable to CR. Thus, the slight modification of the original data set is not evident to the performance of the fab.

Data Sets 5a and 6 We observed similar results for data sets 5a and 6. For this reason and since for data set 5 only very few lots finish their process flow completely we do not show this results explicitly. As for data sets 1 and 3 the improvement in terms of smaller variance of cycle times depends on the fab and on the product. Furthermore, in some cases EDD outperforms FIFO and vice versa.

4 Conclusion and Outlook

In this paper we aimed at reducing the variability of cycle times in semiconductor manufacturing lines by reducing the magnitude of overtaking. The results obtained show that by means of appropriate sequencing rules the amount of overtaking can be reduced. However, contrarily to the original conjecture, this reduction does not necessarily lead to smaller variances of cycle times. The reduction depends very much on the facility under consideration and its corresponding product mix. Thus, no general correlation between the magnitude of overtaking and the variance of cycle times can be established. Nevertheless, our study showed us that the choice of sequencing rules is too important to be ignored. Our findings are potentially meaningful for the design of job floor planning and control systems, since they reinforce the need of advanced methods to evaluate the performance of complex manufacturing systems like semiconductor lines at the fab design level in order to make the right decision which sequencing rule to be employed.

We also have to emphasize that the departure process of the semiconductor manufacturing line is the input process to the backend stage, i.e. a facility where semiconductor products are packed for shipping or used as assembly parts for any kind of machinery. It is a well know fact from queueing theory that highly variable or even correlated input to a queueing system has a tremendous negative effect on its performance. In other words, if the output process of a system suffers from

variability it hurts the performance of downstream production stages. Consequently, further research has do be done to identify and examine the various sources of variance that can affect the variability of cycle times. We also believe that other statistical performance measures than the mean and variance should be observed to characterize wafer cycle times and evaluate semiconductor manufacturing facilities. Additionally, based on these measures, a parameter fitting or spectral analysis of the data may be performed so as to obtain an approximate cycle time distribution.

Acknowledgement The authors would like to thank the MIMAC team for many helpful discussions on data sets and Delphi, in particular Frank Chance, John Fowler, Ottmar Gihr, Luc de Ridder, Jen Robinson, and Ben Rodriguez. The authors would also like to thank Prof. P. Tran-Gia for the support of this work.

References

[1] Arnold O. Allen, 1990. *Probability, Statistics, and Queueing Theory.* Academic Press, London, 2nd edition, 1990.

[2] Kenneth R. Baker, 1984. Sequencing rules and due-date assignments in a job shop. *Management Science*, 30(9):1093–1104, September 1984.

[3] Frank Chance, 1994. *DELPHI: A C–based manufacturing simulator.* Cornell University, Ithaca, NY, U.S.A., 1994. Version 8, Level 29.

[4] Hermann Gold and Bernhard Frötschl, 1994. Performance analysis of a batch service system. In Jacques Labetoulle and James W. Roberts, editors, *Proc. 14th International Teletraffic Congress (ITC 14)*, pages 155–168, Antibes Juan-les-Pins, 1994.

[5] P. R. Kumar, 1993. Re–entrant lines. *Queueing Systems*, 13:87–110, 1993.

[6] Averill M. Law and W. David Kelton, 1991. *Simulation Modeling & Analysis.* McGraw–Hill, New York, 2nd edition, 1991.

[7] Robert C. Leachman and Tali F. Carmon, 1992. On capacity modeling for production planning with alternative machine types. *IIE Transactions*, 24(2):62–72, 1992.

[8] Benjamin Melamed, 1982. Sojourn times in queueing networks. *Mathematics of Operations Research*, 7:223–244, May 1982.

[9] Manfred Mittler and Notker Gerlich, 1993. Reducing the variance of sojourn times in queueing networks with overtaking. Forschungsbericht, Preprint-Reihe Nr. 73, Universität Würzburg, Institut für Informatik, November 1993.

[10] Michael Pinedo, 1995. *Scheduling: Theory, Algorithms, and Systems.* International Series in Industrial and Systems Engineering. Prentice Hall, Englewood Cliffs, 1995.

[11] Lee W. Schruben, H. Singh, and L. Tierney, 1983. Optimal tests for initialization bias in simulation output. *Operations Research*, 31(6):1167–1178, November-December 1983.

[12] Phuoc Tran-Gia, Herman Gold, and Heidrun Grob, 1993. A batch service system operating in a pull production line. *Archiv für Elektronik und Übertragungstechnik*, 47(5/6):379–385, 1993.

[13] Jean Walrand and P. Varaiya, 1980. Sojourn times and the overtaking condition in Jacksonian networks. *Advances in Applied Probability*, 12:1000–1018, 1980.

[14] Ward Whitt, 1984. The amount of overtaking in a network of queues. *Networks*, 14:411–426, 1984.

[15] Susan H. Xu, Rhonda Righter, and J. George Shanthikumar, 1992. Optimal dynamic assignment of customers to heterogeneous servers in parallel. *Operations Research*, 40(6):1126–1138, November-December 1992.